

CLAIMS

WHAT IS CLAIMED IS:

1. A method for servicing, at a predetermined service rate, a plurality of queues containing data packets, each of the queues being associated with respective one of a plurality of connections, each of the connections traversing an associated communication switch, each of the connections being allocated a respective data transfer rate, the method comprising the steps of:
 - responsive to receiving a request for establishing a new connection, determining whether the data transfer rate requested for the new connection matches one of a finite and discrete plurality of basic data transfer rates;
 - responsive to a match of the data transfer rate requested for the new connection with one of the available basic data transfer rates, associating the new connection with the matched basic data transfer rate;
 - responsive to a non-match of the data transfer rate requested for the new connection with any of the available basic data transfer rates, splitting the new connection into a plurality of sub-connections, the allocated rate of each of the sub-connections matching one of the basic data transfer rates, and associating each of the sub-connections with the respective matched basic data transfer rate;
 - responsive to receiving a plurality of data packets via a plurality of data links, identifying for each received data packet the respective one of the allocated connections and identifying the one of the queues associated with the connection;
 - storing each of the received data packets in the corresponding identified queue;
 - responsive to storing each received data packet in a queue associated with one of those connections whose data transfer rate matches one of the available basic data transfer rates,

determining whether the connection was previously idle and must be activated, and identifying the connection as an active connection;

responsive to storing a received data packet in a queue associated with one of those connections whose data transfer rate does not match any of the available basic data transfer rates, in which the connection is identified as a split connection, determining whether a previously idle sub-connection must be activated for the split connection, and identifying the sub-connection as an active sub-connection;

responsive to the activation of a previously idle sub-connection, treating the sub-connection as an independent entity for the purpose of service distribution, the service distribution being regulated by a scheduling method capable of supporting the finite and discrete plurality of basic data transfer rates;

scheduling the transmission of stored data packets on the available outputs according to the data transfer rates associated with respective ones of the allocated connections and sub-connections;

responsive to the transmission of a data packet from a queue associated with one of those connections whose data transfer rate matches one of the available basic data transfer rates to an output, determining whether the queue associated with the connection has at least one more packet waiting therein, in which the connection is identified as an active connection, or whether the queue associated with the connection has no packets waiting therein, in which the connection is identified as an idle connection; and

responsive to the transmission of a data packet from a queue associated with one of those connections whose data transfer rate does not match any of the available basic data transfer rates to an output, updating the state of the packet queue associated with the allocated

split connection, and determining whether one of the currently active sub-connections of the allocated split connection must be deactivated, in which the deactivated sub-connection is identified as an idle sub-connection.

5 2. The method of claim 1, wherein the step of splitting a connection generates a plurality of sub-connections whose data transfer rates are all equal and matching one of the available basic data transfer rates.

 3. The method of claim 1, further comprising the steps of:
10 activating a previously idle sub-connection for a connection whose data transfer rate does not match any of the available basic data transfer rates when a new packet is stored in the queue associated with the split connection and the number of packets currently stored in the queue is smaller than the total number of currently active sub-connections associated with the split connection, in which the activated sub-connection is identified as an active sub-connection;
15 and

 deactivating a previously active sub-connection for a connection whose data transfer rate does not match any of the available basic data transfer rates when a packet is transmitted to an output from the queue associated with the split connection and the number of packets remaining in the queue after the transmission is completed is smaller than the total
20 number of currently active sub-connections associated with the split connection, in which the deactivated sub-connection is identified as an idle sub-connection.

4. The method of claim 1, wherein a macro-queue of connections and sub-connections is associated with each of the plurality of basic data transfer rates, the macro-queue being used to store active connections and active sub-connections whose associated data transfer rate matches the one of the available basic data transfer rates; and

5 wherein the method further comprises the steps of:

appending a connection or sub-connection to the tail of the associated macro-queue when the connection or sub-connection becomes active after being idle, in which a connection is identified as idle when the connection has no data packets waiting for transmission in its respective packet queue, and is identified as active when the packet queue contains at least one data packet, and in which a sub-connection is identified as idle when it has been deactivated, and is identified as active when it has been activated;

associating a timestamp with each of the macro-queues;

generating a new timestamp associated with a macro-queue each time a new active connection or sub-connection reaches the head of the macro-queue, wherein a system potential is used in the generation of the new timestamp;

15 associating a cumulative service rate with each of the macro-queues, the cumulative service rate being used to generate the respective timestamp associated with each of the macro-queues;

generating a value for the system potential according to a predetermined function; and

20 selecting one of the timestamps associated with the macro-queues which have at least one active connection or sub-connection waiting for service therein, and identifying

the active connection or sub-connection at the head of the macro-queue associated with the selected timestamp as the recipient of the next service, the service including the steps of:

removing a data packet from the head of the queue associated with the identified active connection or with the split connection associated with the identified active sub-connection;

transmitting the removed data packet to an output;

removing the identified active connection or sub-connection from the head of the macro-queue associated with the selected timestamp; and

appending again the identified connection or sub-connection to the tail of the same macro-queue only if the identified connection or sub-connection remains active after the packet has been removed.

5. The method of claim 4, wherein the selection of the timestamp identifying an active connection or sub-connection as the recipient of a service is based on a Smallest-

Finishing-potential-First (SFF) timestamp selection policy, wherein a timestamp is selected when its value is minimum among all timestamps associated with non-empty macro-queues.

6. The method of claim 4, wherein the selection of the timestamp identifying an active connection or sub-connection as the recipient of a service is based on a Smallest-Eligible-

Finishing-potential-First (SEFF) timestamp selection policy, wherein a timestamp is selected when its value is minimum among all timestamps being associated with non-empty macro-queues and not exceeding the system potential by more than a fixed amount, the amount being specific of each timestamp.

7. The method of claim 4, wherein the cumulative rate associated with each macro-queue is equal to the basic data transfer rate associated with the macro-queue multiplied by the number of allocated connections and sub-connections associated with the macro-queue.

5

8. The method of claim 4, wherein the timestamp associated with a macro-queue, generated each time a new active connection or sub-connection reaches the head of the macro-queue, is generated as a maximum value between a previous value of the timestamp assigned to the macro-queue and a current value of the system potential, the maximum value incremented by
10 the inverse of the cumulative data transfer rate associated with the macro-queue, if the macro-queue is empty before its head is reached by the new active connection or sub-connection.

9. The method of claim 4, further comprising the steps of:

15 activating a previously idle sub-connection for a connection whose data transfer rate does not match any of the available basic data transfer rates when a new packet is stored in the queue associated with the split connection and the number of packets currently stored in the queue is smaller than the total number of currently active sub-connections associated with the split connection, in which the activated sub-connection is identified as an active sub-connection;
and

20 deactivating a previously active sub-connection for a connection whose data transfer rate does not match any of the available basic data transfer rates when a packet is transmitted to an output from the queue associated with the split connection and the number of packets remaining in the queue after the transmission is completed is smaller than the total

number of currently active sub-connections associated with the split connection, in which the deactivated sub-connection is identified as an idle sub-connection.

10. The method of claim 1, wherein a macro-queue of connections and sub-connections is associated with each of the plurality of basic data transfer rates, the macro-queue being used to store active connections and sub-connections whose associated data transfer rate matches the one of the available basic data transfer rates; and

wherein the method further comprises the steps of:

appending a connection or sub-connection to the tail of the associated macro-queue when the connection or sub-connection becomes active after being idle, in which a connection is identified as idle when it has no data packets waiting for transmission in its respective packet queue, and is identified as active when the packet queue contains at least one data packet, and in which a sub-connection is identified as idle when it has been deactivated, and is identified as active when it has been activated;

associating a timestamp with each of the active connections and sub-connections;

generating a new timestamp associated with one of the active connections and sub-connections each time the active connection or sub-connection is appended to the tail of the associated macro-queue, wherein a system potential is used in the generation of the new timestamp;

associating a timestamp with each macro-queue that contains at least one active connection or sub-connection, the timestamp being the one associated with the active connection or sub-connection currently located at the head of the macro-queue;

using the basic data transfer rate associated with each of the macro-queues
to generate the timestamps associated with the respective connections and sub-connections;
generating a value for the system potential according to a predetermined
function; and

5 selecting one of the timestamps associated with those of the macro-queues
which have at least one active connection or sub-connection waiting for service therein, and
identifying the active connection or sub-connection at the head of the macro-queue associated
with the selected timestamp as the recipient of the next service, the service including the steps of:

removing a data packet from the head of the queue associated with
10 the identified active connection or with the split connection associated with the
identified active sub-connection;

transmitting the removed data packet to an output;

removing the identified active connection or sub-connection from
the head of the macro-queue associated with the selected timestamp; and

15 appending again the identified connection or sub-connection to the
tail of the same macro-queue only if the identified connection or sub-connections
remains active after the packet has been removed.

11. The method of claim 10, wherein the selection of the timestamp identifying an
20 active connection or sub-connection as the recipient of a service is based on a Smallest-
Finishing-potential-First (SFF) timestamp selection policy, wherein a timestamp is selected when
its value is minimum among all timestamps associated with non-empty macro-queues.

12. The method of claim 10, wherein the selection of the timestamp identifying an active connection or sub-connection as the recipient of a service is based on a Smallest-Eligible-Finishing-potential-First (SEFF) timestamp selection policy, wherein a timestamp is selected when its value is minimum among all timestamps being associated with non-empty macro-
5 queues and not exceeding the system potential by more than a fixed amount, the amount being specific of each timestamp.

13. The method of claim 10, wherein the timestamp associated with a connection or sub-connection, generated each time the connection or sub-connection is queued to the respective
10 macro-queue, is generated as a maximum value between a previous value of the timestamp assigned to the connection or sub-connection and a current value of the system potential, the maximum value incremented by the inverse of the data transfer rate associated with the macro-queue, if the connection or sub-connection is not active before being queued to the macro-queue.

14. The method of claim 10, further comprising the steps of:
activating a previously idle sub-connection for a connection whose data transfer
rate does not match any of the available basic data transfer rates when a new packet is stored in
the queue associated with the split connection and the number of packets currently stored in the
queue is smaller than the total number of currently active sub-connections associated with the
20 split connection, in which the activated sub-connection is identified as an active sub-connection;
and

deactivating a previously active sub-connection for a connection whose data
transfer rate does not match any of the available basic data transfer rates when a packet is

transmitted to an output from the queue associated with the split connection and the number of packets remaining in the queue after the transmission is completed is smaller than the total number of currently active sub-connections associated with the split connection, in which the deactivated sub-connection is identified as an idle sub-connection.

5

15. A system for servicing, at a predetermined service rate, a plurality of queues containing data packets, each of the queues being associated with respective one of a plurality of connections, each of the connections traversing an associated communication switch, each of the connections being allocated a respective data transfer rate, the system comprising:

10 means, responsive to receiving a request for establishing a new connection, for determining whether the data transfer rate requested for the new connection matches one of a finite and discrete plurality of basic data transfer rates;

means, responsive to a match of the data transfer rate requested for the new connection with one of the available basic data transfer rates, for associating the new connection
15 with the matched basic data transfer rate;

means, responsive to a non-match of the data transfer rate requested for the new connection with any of the available basic data transfer rates, for splitting the new connection into a plurality of sub-connections, the data transfer rate of each of the sub-connections matching one of the basic data transfer rates, and associating each of the sub-connections with the
20 respective matched basic data transfer rate;

means, responsive to receiving a plurality of data packets via a plurality of data links, for identifying for each received data packet the respective one of the allocated connections and identifying the one of the queues associated with the connection;

means for storing each of the received data packets in the corresponding identified queue;

means, responsive to storing each received data packet in a queue associated with one of those connections whose data transfer rate matches one of the available basic data transfer rates, for determining whether the connection was previously idle and must be activated, and for
5 identifying the connection as an active connection;

means, responsive to storing a received data packet in a queue associated with one of those connections whose data transfer rate does not match any of the available basic data transfer rates, in which the connection is identified as a split connection, for determining whether
10 a previously idle sub-connection must be activated for the split connection, and for identifying the activated sub-connection as an active sub-connection;

means, responsive to the activation of a previously idle sub-connection, for treating the sub-connection as an independent entity for the purpose of service distribution, the service distribution being regulated by a scheduling method capable of supporting the finite and
15 discrete plurality of basic data transfer rates;

means for scheduling the transmission of stored data packets on the available outputs according to the data transfer rates associated with respective ones of the allocated connections and sub-connections;

means, responsive to the transmission of a data packet from a queue associated
20 with one of those connections whose data transfer rate matches one of the available basic data transfer rates to an output, for determining whether the queue associated with the connection has at least one more packet waiting therein, in which the connection is identified as an active

connection, or whether the queue associated with the connection has no packets waiting therein, in which the connection is identified as an idle connection; and

means, responsive to the transmission of a data packet from a queue associated with one of those connections whose data transfer rate does not match any of the available basic data transfer rates to an output, for updating the state of the packet queue associated with the allocated split connection, and for determining whether one of the currently active sub-connections of the allocated split connection must be deactivated, in which the sub-connection is identified as an idle sub-connection.

16. The system of claim 15, wherein the means for splitting a connection generates a plurality of sub-connections whose data transfer rates are all equal and matching one of the available basic data transfer rates.

17. The system of claim 15, further comprising:

means for activating a previously idle sub-connection for a connection whose data transfer rate does not match any of the available basic data transfer rates when a new packet is stored in the queue associated with the split connection and the number of packets currently stored in the queue is smaller than the total number of currently active sub-connections associated with the split connection, in which the activated sub-connection is identified as an active sub-connection; and

means for deactivating a previously active sub-connection for a connection whose data transfer rate does not match any of the available basic data transfer rates when a packet is transmitted to an output from the queue associated with the split connection and the number of

packets remaining in the queue after the transmission is completed is smaller than the total number of currently active sub-connections associated with the split connection, in which the deactivated sub-connection is identified as an idle sub-connection.

5 18. The system of claim 15, further comprising:

 a macro-queue for storing connections and sub-connections associated with each of the plurality of basic data transfer rates, the macro-queue being used to store active connections and sub-connections whose associated data transfer rate matches the one of the available basic data transfer rates;

10 means for appending a connection or sub-connection to the tail of the associated macro-queue when the connection or sub-connection becomes active after being idle, in which a connection is identified as idle when the connection has no data packets waiting for transmission in its respective packet queue, and is identified as active when the packet queue contains at least one data packet, and in which a sub-connection is identified as idle when it has been deactivated,
15 and is identified as active when it has been activated;

 means for associating a timestamp with each of the macro-queues;

 means for generating a new timestamp associated with a macro-queue each time a new active connection or sub-connection reaches the head of the macro-queue, wherein a system potential is used in the generation of the new timestamp;

20 means for associating a cumulative service rate with each of the macro-queues, the cumulative service rate being used to generate the respective timestamp associated with each of the macro-queues;

means for generating a value for the system potential according to a predetermined function; and

means for selecting one of the timestamps associated with the macro-queues which have at least one active connection or sub-connection waiting for service therein, and for
5 identifying the active connection or sub-connection at the head of the macro-queue associated with the selected timestamp as the recipient of the next service, the service including the steps of:

removing a data packet from the head of the queue associated with the identified active connection or with the split connection associated with the identified active sub-connection;

10 transmitting the removed data packet to an output;

removing the identified active connection or sub-connection from the head of the macro-queue associated with the selected timestamp; and

appending again the identified connection or sub-connection to the tail of the same macro-queue only if the identified connection or sub-connection remains
15 active after the packet has been removed.

19. The system of claim 18, wherein the means for selecting the timestamp that identifies an active connection or sub-connection as the recipient of a service is based on a Smallest-Finishing-potential-First (SFF) timestamp selection policy, wherein a timestamp is
20 selected when its value is minimum among all timestamps associated with non-empty macro-queues.

20. The system of claim 18, wherein the means for selecting the timestamp that identifies an active connection or sub-connection as the recipient of a service is based on a Smallest-Eligible-Finishing-potential-First (SEFF) timestamp selection policy, wherein a timestamp is selected when its value is minimum among all timestamps being associated with non-empty macro-queues and not exceeding the system potential by more than a fixed amount, the amount being specific of each timestamp.

21. The system of claim 18, wherein the cumulative rate associated with each macro-queue is equal to the basic data transfer rate associated with the macro-queue multiplied by the number of allocated connections and sub-connections associated with the macro-queue.

22. The system of claim 18, wherein the timestamp associated with a macro-queue, generated each time a new active connection or sub-connection reaches the head of the macro-queue, is generated as a maximum value between a previous value of the timestamp assigned to the macro-queue and a current value of the system potential, the maximum value incremented by the inverse of the cumulative data transfer rate associated with the macro-queue, if the macro-queue is empty before its head is reached by the new active connection or sub-connection.

23. The system of claim 18, further comprising:
means for activating a previously idle sub-connection for a connection whose data transfer rate does not match any of the available basic data transfer rates when a new packet is stored in the queue associated with the split connection and the number of packets currently stored in the queue is smaller than the total number of currently active sub-connections

associated with the split connection, in which the activated sub-connection is identified as an active sub-connection; and

means for deactivating a previously active sub-connection for a connection whose data transfer rate does not match any of the available basic data transfer rates when a packet is transmitted to an output from the queue associated with the split connection and the number of packets remaining in the queue after the transmission is completed is smaller than the total number of currently active sub-connections associated with the split connection, in which the deactivated sub-connection is identified as an idle sub-connection.

24. The system of claim 15, wherein a macro-queue of connections and sub-connections is associated with each of the plurality of basic data transfer rates, the macro-queue being used to store active connections and sub-connections whose associated data transfer rate matches the one of the available basic data transfer rates; and

wherein the system further comprises:

means for appending a connection or sub-connection to the tail of the associated macro-queue when the connection or sub-connection becomes active after being idle, in which a connection is identified as idle when it has no data packets waiting for transmission in its respective packet queue, and is identified as active when the packet queue contains at least one data packet, and in which a sub-connection is identified as idle when it has been deactivated, and is identified as active when it has been activated;

means for associating a timestamp with each of the active connections and sub-connections;

means for generating a new timestamp associated with one of the active connections and sub-connections each time the active connection or sub-connection is appended to the tail of the associated macro-queue, wherein a system potential is used in the generation of the new timestamp;

5 means for associating a timestamp with each macro-queue that contains at least one active connection or sub-connection, the timestamp being the one associated with the active connection or sub-connection currently located at the head of the macro-queue;

means for using the basic data transfer rate associated with each of the macro-queues to generate the timestamps associated with the respective connections and sub-connections;

10

means for generating a value for the system potential according to a predetermined function; and

means for selecting one of the timestamps associated with those of the macro-queues which have at least one active connection or sub-connection waiting for service therein, and identifying the active connection or sub-connection at the head of the macro-queue associated with the selected timestamp as the recipient of the next service, the service including the steps of:

15

removing a data packet from the head of the queue associated with the identified active connection or with the split connection associated with the identified active sub-connection;

20

transmitting the removed data packet to an output;

removing the identified active connection or sub-connection from the head of the macro-queue associated with the selected timestamp; and

appending again the identified connection or sub-connection to the tail of the same macro-queue only if the identified connection or sub-connection remains active after the packet has been removed.

5 25. The system of claim 24, wherein the selection of the timestamp identifying an active connection or sub-connection as the recipient of a service is based on a Smallest-Finishing-potential-First (SFF) timestamp selection policy, wherein a timestamp is selected when its value is minimum among all timestamps associated with non-empty macro-queues.

10 26. The system of claim 24, wherein the selection of the timestamp identifying an active connection or sub-connection as the recipient of a service is based on a Smallest-Eligible-Finishing-potential-First (SEFF) timestamp selection policy, wherein a timestamp is selected when its value is minimum among all timestamps being associated with non-empty macro-queues and not exceeding the system potential by more than a fixed amount, the amount being
15 specific of each timestamp.

 27. The system of claim 24, wherein the timestamp associated with a connection or sub-connection, generated each time the connection or sub-connection is queued to the respective macro-queue, is generated as a maximum value between a previous value of the timestamp
20 assigned to the connection or sub-connection and a current value of the system potential, the maximum value incremented by the inverse of the data transfer rate associated with the macro-queue, if the connection or sub-connection is not active before being queued to the macro-queue.

28. The system of claim 24, further comprising:

means for activating a previously idle sub-connection for a connection whose data transfer rate does not match any of the available basic data transfer rates when a new packet is stored in the queue associated with the split connection and the number of packets currently stored in the queue is smaller than the total number of currently active sub-connections associated with the split connection, in which the activated sub-connection is identified as an active sub-connection; and

means for deactivating a previously active sub-connection for a connection whose data transfer rate does not match any of the available basic data transfer rates when a packet is transmitted to an output from the queue associated with the split connection and the number of packets remaining in the queue after the transmission is completed is smaller than the total number of currently active sub-connections associated with the split connection, in which the deactivated sub-connection is identified as an idle sub-connection.

29. A scheduler for processing a plurality of connections in a packet-processing system, each of the connections having an associated data transfer rate, the scheduler comprising:

a plurality of rate first-in-first-out (FIFO) queues for storing the connections;

a packet selector for servicing the connections; and

a connection controller for performing the steps of:

responsive to receiving a request for establishing a new connection,
determining whether the data transfer rate requested for the new connection matches one of a finite and discrete plurality of basic data transfer rates;

responsive to a match of the data transfer rate requested for the new connection with one of the available basic data transfer rates, associating the new connection with the matched basic data transfer rate;

5 responsive to a non-match of the data transfer rate requested for the new connection with any of the available basic data transfer rates, splitting the new connection into a plurality of sub-connections, the allocated rate of each of the sub-connections matching one of the basic data transfer rates, and associating each of the sub-connections with the respective matched basic data transfer rate;

10 responsive to receiving a plurality of data packets via a plurality of data links, identifying for each received data packet the respective one of the allocated connections and identifying the one of the queues associated with the connection;

storing each of the received data packets in the respective identified queue;

15 responsive to storing each received data packet in a queue associated with one of those connections whose data transfer rate matches one of the available basic data transfer rates, determining whether the connection was previously idle and must be activated, and identifying the connection as an active connection;

20 responsive to storing a received data packet in a queue associated with one of those connections whose data transfer rate does not match any of the available basic data transfer rates, in which the connection is identified as a split connection, determining whether a previously idle sub-connection must be activated for the split connection, and identifying the sub-connection as an active sub-connection;

responsive to the activation of a previously idle sub-connection, treating the sub-connection as an independent entity for the purpose of service distribution, the service distribution being regulated by the a scheduling method capable of supporting the finite and discrete plurality of basic data transfer rates;

5 scheduling the transmission of stored data packets on the available outputs according to the data transfer rates associated with respective ones of the allocated connections and sub-connections;

10 responsive to the transmission of a data packet from a queue associated with one of those connections whose data transfer rate matches one of the available basic data transfer rates to an output, determining whether the queue associated with the connection has at least one more packet waiting therein, in which the connection is identified as an active connection, or whether the queue associated with the connection has no packets waiting therein, in which the connection is identified as an idle connection; and

15 responsive to the transmission of a data packet from a queue associated with one of those connections whose data transfer rate does not match any of the available basic data transfer rates to an output, updating the state of the packet queue associated with the allocated split connection, and determining whether one of the currently active sub-connections of the allocated split connection must be
20 deactivated, in which the sub-connection is identified as an idle sub-connection.

30. The scheduler of claim 29, wherein the packet selector processes the queued connections and sub-connections using a per-connection-timestamp procedure.

31. The scheduler of claim 29, wherein the packet selector processes the queued connections and sub-connections using a no-per-connection-timestamp procedure.

5 32. A method for processing a session in a packet-processing system, the session having a session rate, the method comprising the steps of:

determining whether the session rate matches one of a plurality of basic rates, with each basic rate associated with a respective one of a plurality of rate-specific queues; and

10 splitting the session, in response to a non-match of the session rate with any of the basic rates, into subsessions for queuing into at least one of the rate-specific queues.

33. The method of claim 32 further comprising the step of:

15 queuing the session in an unsplit state, in response to a match of the session rate with any of the basic rates, into a corresponding rate-specific queue having the matching basic rate.

34. The method of claim 32, wherein the step of splitting includes the step of splitting the session into subsessions wherein every subsession has an identical subsession rate.

20 35. The method of claim 34, wherein the subsession rate matches at least one of the plurality of basic rates.

36. The method of claim 35, further comprising the step of:
queuing the subsessions into a corresponding rate-specific queue having an
associated basic rate matching the subsession rate.

5 37. The method of claim 32, wherein the steps of determining and splitting are
performed by a queue controller.

38. The method of claim 32, further comprising the step of:
processing the session and subsessions using a per-connection-timestamp
10 procedure.

39. The method of claim 32, further comprising the step of:
processing the session and subsessions using a no-per-connection-timestamp
procedure.

15 40. The method of claim 32, wherein the rate-specific queues are rate first-in-first-out
(FIFO) queues.

41. A method for processing a session in a packet-processing system, the session
20 having a session rate, the method comprising the steps of:
determining whether the session rate matches one of a plurality of basic rates,
with each basic rate associated with a respective one of a plurality of rate first-in-first-out (FIFO)
queues;

maintaining the session in an unsplit state, in response to a match of the session rate with any of the basic rates; and

splitting the session into subsessions, in response to a non-match of the session rate with any of the basic rates wherein a subsession rate of each subsession matches at least one of the plurality of basic rates.

42. The method of claim 41 further comprising the steps of:

queuing the session in an unsplit state into at least one of the rate FIFO queues having the corresponding matching basic rate, in response to the step of maintaining; and

queuing the subsessions into at least of the rate FIFO queues having the corresponding basic rate matching the subsession rate, in response to the step of splitting.

43. The method of claim 41, wherein the step of splitting includes the step of splitting the session into subsessions wherein every subsession has an identical subsession rate.

44. The method of claim 41, wherein the steps of determining, maintaining, and splitting are performed by a queue controller.

45. The method of claim 42, further comprising the step of:
processing the queued session and subsessions using a per-connection-timestamp procedure.

46. The method of claim 42, further comprising the step of:
processing the queued session and subsessions using a no-per-connection-
timestamp procedure.

5 47. A scheduler for processing a session in a packet-processing system, the session
having a session rate, the scheduler comprising:
a plurality of rate first-in-first-out (FIFO) queues for storing the session, wherein
each FIFO queue is associated with a respective one of a plurality of basic rates;
a packet selector for servicing the session; and
10 a queue controller for determining whether the session rate matches one of the
plurality of basic rates, for maintaining the session in an unsplit state in response to a match of
the session rate with any of the basic rates, and for splitting the session into subsessions in
response to a non-match of the session rate with any of the basic rates wherein a subsession rate
of each subsession matches at least one of the plurality of basic rates.

15 48. The scheduler of claim 47 wherein the queue controller queues the session in an
unsplit state into at least one of the rate FIFO queues having the corresponding matching basic
rate, and queues the subsessions into at least one of the rate FIFO queues having the
corresponding basic rate matching the subsession rate.

20 49. The scheduler of claim 47, wherein every subsession has an identical subsession
rate.

50. The scheduler of claim 47, wherein the packet selector processes the queued session and subsessions using a per-connection-timestamp procedure.

51. The scheduler of claim 47, wherein the packet selector processes the queued
5 session and subsessions using a no-per-connection-timestamp procedure.